

Non-characterizability in belief revision: an application of finite model theory

György Turán*

University of Illinois at Chicago
MTA-SZTE Research Group on Artificial Intelligence

Jon Yaggie*

University of Illinois at Chicago

April 1, 2014

Abstract

A formal framework is given for the characterizability of a class of belief revision operators, defined using minimization over a class of partial preorders, by postulates. It is shown that for partial orders characterizability implies a definability property of the class of partial orders in monadic second-order logic. Based on a non-definability result for a class of partial orders, an example is given of a non-characterizable class of revision operators. This appears to be the first non-characterizability result in belief revision.

1 Introduction

The main approach to belief change is the AGM approach pioneered by [1]. It provides many characterization results for belief change operators in terms of rationality postulates [7]. *Are there cases where no characterization can be given?* Answering this non-axiomatizability question presupposes a formal definition of a postulate. However, as noted in the survey paper [4]

“theories of belief change developed in the AGM tradition are not logics in a strict sense, but rather informal axiomatic theories of belief change. Instead of characterizing the models of belief and belief change in a formalized object language, the AGM approach uses a natural language (ordinary mathematical English) to characterize the mathematical structures under study.”

*Partially supported by NSF grant CCF-0916708

Fermé and Hansson then proceed to describe modal and dynamic logic approaches to modeling belief change (see, e.g., [10]). As far as we know, the question of non-characterizability by postulates has not been considered before in those frameworks either.

In this note we provide a formal framework for studying characterizability, based on the approach of Katsuno and Mendelzon [8]. A revision operator $*$ in [8] is considered to assign a revised knowledge base $K * \varphi$ to every knowledge base K and every revising formula φ . However, the results remain valid if one considers $*$ to act on a fixed knowledge base K and an arbitrary revising formula φ . We are not discussing iterated revision here, and so there is no interaction between the revisions of different knowledge bases. Katsuno and Mendelzon prove the following results.

Theorem 1.1. [8]

a) *There is a finite set of postulates such that a revision operator satisfies these postulates iff there is a faithful total preorder representing it with minimization.*

b) *There is a finite set of postulates such that a revision operator satisfies these postulates iff there is a faithful partial preorder representing it with minimization.*

c) *There is a finite set of postulates such that a revision operator satisfies these postulates iff there is a faithful partial order representing it with minimization.*

Part a) is a finite version of Grove's characterization of the AGM postulates in terms of systems of spheres [5]. The postulates for parts b) and c) are the same, and different from a).

We consider the following general question.

Problem 1. *Let \mathcal{R} be a family of partial preorders. Is there a finite set of postulates such that a revision operator satisfies these postulates iff there is a faithful partial preorder from \mathcal{R} representing it with minimization?*

Our goal is to prove a negative answer for a particular family \mathcal{R} . The formal definition of characterizability aims at providing a formal framework for proving this negative result. Formulating frameworks for other types of non-characterizability results appears to be an interesting topic for future work.

Non-characterizability is proved using a translation from postulates to universal monadic second-order formulas over the language of partial preorders. It follows from this translation that for partial orders postulate characterization of the class of revision operators implies universal monadic second-order definability of the class of partial orders considered. Thus, non-definability of the class of partial orders implies non-characterizability by postulates. Non-definability in monadic second-order logic is a well-studied topic in finite model theory [2, 9]. We give such a non-definability result for a particular class of partial orders. This class is constructed to give a first example of non-characterizability. It remains an interesting problem to find more natural examples. A candidate is the class of 2-dimensional partial orders.

2 Preliminaries

We consider propositional logic knowledge bases K over a fixed finite set of variables. We write K_n to indicate that K is over n variables. Truth assignments (or interpretations) are assignments of truth values to the variables. The set of truth assignments satisfying a formula φ is denoted by $|\varphi|$. Given a set A of truth assignments, $\langle A \rangle$ is some formula φ such that $|\varphi| = A$. A knowledge base is represented by a single formula¹.

Given a knowledge base K , a belief revision operator $*$ assigns a formula $K * \varphi$ to every formula φ . Here φ is called the revising formula, and $K * \varphi$ is called the revised knowledge base.

A partial preorder is $R = (U, \leq)$, where U is a finite ground set and \leq is a reflexive, transitive binary relation. A partial order is, in addition, antisymmetric. We write $a \sim b$ if a and b are incomparable. The comparability graph of R is the undirected graph over U such that for any pair of vertices (a, b) is an edge iff $a \leq b$ or $b \leq a$. An element a is minimal if there is no b such that $b < a$, where $b < a$ iff $b \leq a$ but $a \not\leq b$. If $U' \subseteq U$ then a is minimal in U' if $a \in U'$ and there is no $b \in U'$ such that $b < a$. The set of minimal elements of U' is denoted by $\min_{\leq} U'$.

Definition 2.1. (*Faithful partial preorder*) A faithful partial preorder for a knowledge base K_n is a pair $F = (R, t)$, where $R = (U, \leq)$ is a partial preorder on 2^n elements and $t : U \rightarrow \{0, 1\}^n$ is a bijection between the elements of U and truth assignments, such that

1. $a \in U$ is minimal iff $t(a)$ satisfies K_n ,
2. if $t(a)$ satisfies K_n and $t(b)$ falsifies K_n then $a < b$.

In the standard definition the partial preorder is defined over the set of truth assignments. For our discussion it is more convenient to separate the partial preorder and the labeling of its elements by truth assignments.

The basic construction used in Theorem 1.1 is that of a revision operator determined by a faithful partial preorder using minimization.

Definition 2.2. (*Revision using minimization*) The revision operator $*_F$ for K , determined by a faithful partial preorder F for K , using minimization is

$$K *_F \varphi = \langle \min_{\leq} t^{-1}(|\varphi|) \rangle.$$

Thus the revised knowledge base is satisfied by the minimal satisfying truth assignments of the revising formula. Faithfulness implies that if the revising formula is consistent with the knowledge base then the revised knowledge base is the conjunction of the knowledge base and the revising formula.

We will use some notions from finite model theory. General introductions to the topic are given in [2, 9]. The notions used are introduced in the later sections, so our discussion is essentially self-contained.

¹We note that because of finiteness this representation corresponds to the belief set framework. Computational complexity issues are not discussed here thus the details of the representation are irrelevant.

3 Postulates

Consider the AGM postulates

$$\text{if } K \text{ is satisfiable then } K * \varphi \text{ is also satisfiable} \quad (1)$$

and

$$\text{if } (K * \varphi) \wedge \psi \text{ is satisfiable then } K * (\varphi \wedge \psi) \vdash (K * \varphi) \wedge \psi. \quad (2)$$

Here $K, K * \varphi, \varphi$ and ψ can be considered as unary predicates over the set of interpretations, and thus (1) can be rewritten as

$$[\exists x K(x)] \rightarrow [\exists x (K * \varphi)(x)]$$

and (2) can be rewritten as

$$[\exists x ((K * \varphi)(x) \wedge \psi(x))] \rightarrow [\forall y ((K * (\varphi \wedge \psi))(y) \rightarrow ((K * \varphi)(y) \wedge \psi(y))]. \quad (3)$$

Postulates refer to a fixed knowledge base K , and are implicitly universally quantified over formula symbols such as φ, ψ . They express general requirements that are supposed to hold for all revising formulas. Generalizing these examples, a postulate is defined as follows.

Definition 3.1. (*Postulate*) A postulate P is a first-order sentence with unary predicate symbols $K, \varphi_1, \dots, \varphi_\ell$ and $K * \mu_1, \dots, K * \mu_m$, where μ_1, \dots, μ_m are Boolean combinations of $\varphi_1, \dots, \varphi_\ell$.

A revision operator satisfies a postulate for a knowledge base K if the postulate holds for all $\varphi_1, \dots, \varphi_\ell$, with the variables ranging over the set of truth assignments.

This definition covers all postulates in [8] and in Section 7.3 of [7].

4 Characterizability

As we consider partial preorders that are faithful for a knowledge base, we introduce the following property of partial preorders.

Definition 4.1. (*Regular partial preorders*) A partial preorder is regular if

1. every minimal element is smaller than any non-minimal element
2. the number of elements is a power of 2.

An example of a non-regular partial preorder is the 4-element partial order with $a < b, c < d$ and no other comparability. Condition 1 is satisfied, for example, if there is a unique minimal element.

Definition 4.2. (*R-revision operator*) Let \mathcal{R} be a family of regular partial preorders. Let K be a knowledge base and $*$ be a revision operator for K . Then $*$ is an \mathcal{R} -revision operator iff there is a faithful partial preorder $F = (R, t)$ for K , with $R \in \mathcal{R}$, representing $*$ using minimization.

Using this definition, a formal definition is given of characterizability.

Definition 4.3. (*Characterization, characterizability*) A finite set of postulates \mathcal{P} characterizes \mathcal{R} -revision operators if for every knowledge base K and every revision operator $*$ for K the following holds: $*$ satisfies the postulates in \mathcal{P} iff it is an \mathcal{R} -revision operator.

The family of \mathcal{R} -revision operators is characterizable if there is a finite set of postulates characterizing \mathcal{R} -revision operators.

It may be assumed *w.l.o.g.* that \mathcal{P} consists of a single postulate.

5 A non-characterizable class

A partial order is a *crown* if it has elements a_1, \dots, a_s and b_1, \dots, b_s for some s , and the comparabilities are $a_i > b_i$ and $a_i > b_{i+1}$ for every i , where the indices are meant cyclically. A partial order is a *double crown* if it consists of two crowns with pairwise incomparable elements. An *extended double crown* has additional elements that are smaller than all other elements. An *extended crown* is a crown with additional elements that are smaller than all other elements. Thus extended double crowns and extended crowns satisfy the first condition of Definition 4.1 and if their size is a power of 2 then they are regular.

Let \mathcal{R}_0 be the family of extended double crowns with size a power of 2, and \mathcal{R}_1 be the family of extended crowns with size a power of 2. The following theorem gives a sufficient condition for non-characterizability. Its proof is given in the next two sections.

Theorem 5.1. *Let \mathcal{R} be any family of regular partial orders containing \mathcal{R}_0 and disjoint from \mathcal{R}_1 . Then the family of \mathcal{R} -revision operators is not characterizable.*

As a special case of the theorem, we formulate one specific non-characterizable family of revision operators. A partial preorder R is *regular-disconnected* if it is regular, and the partial preorder obtained from R by removing its minimal elements has a disconnected comparability graph.

Corollary 5.2. *Let \mathcal{R} be the family of regular-disconnected partial orders. Then the family of \mathcal{R} -revision operators is not characterizable.*

The corollary follows directly from Theorem 5.1 and the definitions.

6 Translation lemma

Now we define a translation of postulates into sentences over an extension of the language of partial preorders. The language of partial preorders contains a binary relation symbol \leq and equality.

The translated sentences also contain additional unary predicate symbols A_1, \dots, A_ℓ . These correspond to propositional formulas $\varphi_1, \dots, \varphi_\ell$ occurring in the postulates. Given a Boolean combination μ of $\varphi_1, \dots, \varphi_\ell$, we denote by $\hat{\mu}$

the first-order formula obtained by replacing the φ 's with A 's. For instance, for $\mu(x) = (\varphi_1 \wedge \varphi_2)(x)$ becomes $\hat{\mu}(x) = A_1(x) \wedge A_2(x)$.

Given a formula ν over the language \leq, A_1, \dots, A_ℓ with a single free variable x we write \min_{\leq}^ν for a formula expressing that x is a minimal element satisfying ν , i.e.,

$$\min_{\leq}^\nu(x) \equiv \nu(x) \wedge \forall y(\nu(y) \rightarrow \neg(y < x)).$$

Minimal elements in the partial preorder are defined by

$$\min_{\leq}(x) \equiv \forall y(\neg(y < x)).$$

Definition 6.1. (*Translation*) The translation $\tau(P)$ of a postulate P is the sentence obtained from P by replacing

1. every occurrence of $K(x)$ with $\min_{\leq}(x)$
2. every occurrence of $\varphi_i(x)$ and $\mu_i(x)$ with their “hat” versions
3. every occurrence of $K * \mu_i$ with $\min_{\leq}^{\mu_i}(x)$

The translation is a first-order sentence over the predicate symbols \leq, A_1, \dots, A_ℓ .

Example 6.2. (*Translation of postulate (3)*) Let us replace φ and ψ with φ_1 and φ_2 to be consistent with the general notation:

$$[\exists x((K * \varphi_1)(x) \wedge \varphi_2(x))] \rightarrow [\forall y((K * (\varphi_1 \wedge \varphi_2))(y) \rightarrow ((K * \varphi_1)(y) \wedge \varphi_2(y))].$$

Applying Definition 6.1 we get

$$[\exists x(\min_{\leq}^{A_1}(x) \wedge A_2(x))] \rightarrow [\forall y(\min_{\leq}^{A_1 \wedge A_2}(y) \rightarrow (\min_{\leq}^{A_1}(y) \wedge A_2(y))].$$

Given $K, \varphi_1, \dots, \varphi_k$ and a faithful partial preorder F for K , the $(\varphi_1, \dots, \varphi_k)$ -extension of F is determined in the standard way, by interpreting the unary predicate symbols A_1, \dots, A_k by $A_i(a) = \varphi_i(t(a))$. The following proposition is a direct consequence of the definitions.

Proposition 6.3. Let K be a knowledge base, $F = (R, t)$ be a faithful partial preorder for K and let $*_F$ be the revision operator determined by F using minimization. Let $\varphi_1, \dots, \varphi_\ell$ be propositional formulas and P be a postulate. Then P is satisfied by K for $\varphi_1, \dots, \varphi_\ell$ iff the $(\varphi_1, \dots, \varphi_\ell)$ -extension of F satisfies $\tau(P)$.

Now we formulate the connection to definability. In the formulation of we restrict ourselves to partial orders.

A universal monadic second-order sentence is of the form

$$\Phi = \forall A_1, \dots, A_\ell \Psi,$$

where A_1, \dots, A_ℓ range over unary predicates (or subsets) over the universe, and Ψ is a first-order sentence using the unary predicate symbols A_1, \dots, A_ℓ in addition to the original language (in our case \leq and equality). An existential second-order sentence is of the form $\Phi = \exists A_1, \dots, A_\ell \Psi$.

Definition 6.4. (*Universal monadic second-order definability*) A family \mathcal{R} of regular partial orders is universal monadic second-order definable if there is a universal monadic second-order sentence Φ such that for every regular partial order R it holds that $R \in \mathcal{R}$ iff it satisfies Φ .

Lemma 6.5. Let \mathcal{R} be a family of regular partial preorders. If the family of \mathcal{R} -revision operators is characterizable then \mathcal{R} is universal monadic second-order definable.

Proof Let \mathcal{R} be a family of regular partial orders such that \mathcal{R} -revision operators are characterized by a postulate P . We claim that \mathcal{R} is defined by the universal monadic second-order sentence

$$\Phi = \forall A_1, \dots, A_\ell \tau(P).$$

Assume that the regular partial order $R = (U, \leq)$ is in \mathcal{R} . Let the number of its elements be 2^n . Let $t : U \rightarrow \{0, 1\}^n$ be an arbitrary bijection between U and the set of truth assignments. We get a faithful partial preorder $F = (R, t)$ for some knowledge base K_n , and thus the corresponding revision operator $*_F$ is an \mathcal{R} -revision operator. Therefore $*_F$ satisfies P . Consider arbitrary unary relations A_1, \dots, A_ℓ over the elements. Applying Proposition 6.3 to the propositional formulas $\varphi_1, \dots, \varphi_\ell$ corresponding to A_1, \dots, A_ℓ , it follows that \leq, A_1, \dots, A_ℓ satisfy $\tau(P)$. Thus R satisfies Φ .

Now assume that the regular partial order R is not in \mathcal{R} . Again, let $t : U \rightarrow \{0, 1\}^n$ be an arbitrary bijection between U and the set of truth assignments. We get a faithful partial preorder $F = (R, t)$ for some knowledge base K_n . This determines a revision operator $*_F$.

We claim that $*_F$ is not an \mathcal{R} -revision operator. This follows if we show that F is the only faithful partial order determining $*_F$. Assume that for $F' = (R', t')$ with $R' = (U', \leq')$ the revision operator $*_{F'}$ is the same. Then, as revision operators are defined using minimization, for every pair of truth assignments u, v it holds that

$$t^{-1}(u) < t^{-1}(v) \text{ iff } K_n *_F \langle u, v \rangle = \langle u \rangle \text{ iff } (t')^{-1}(u) < (t')^{-1}(v)$$

and

$$t^{-1}(u) \sim t^{-1}(v) \text{ iff } K_n *_F \langle u, v \rangle = \langle u, v \rangle \text{ iff } (t')^{-1}(u) \sim (t')^{-1}(v).$$

Thus $(t')^{-1} \circ t$ is an isomorphism from R to R' .

Hence $*_F$ does not satisfy P . So there are propositional formulas $\varphi_1, \dots, \varphi_\ell$ such that the corresponding instance of P is false. By Proposition 6.3 the corresponding unary predicates A_1, \dots, A_ℓ falsify $\tau(P)$. Hence R falsifies Φ . \square

7 Proof of Theorem 5.1

Two first-order structures are q -equivalent (denoted by \equiv_q) if they satisfy the same first-order sentences of quantifier rank at most q . The q -round Ehrenfeucht - Fraissé game over two structures is played by two players, Spoiler and

Duplicator. In each round Spoiler picks one of the structures and an element of that structure. Duplicator responds by picking an element in the other structure. After q rounds Duplicator wins if the substructures of picked elements in the two structures are isomorphic. Otherwise Spoiler wins. Duplicator has a winning strategy iff the two structures are q -equivalent.

The following result uses the notion of neighborhood in a general relational structure. We only use this result for undirected graphs with colored vertices, where the r -neighborhood of a vertex v is the set of vertices reachable from v by paths of length at most r .

Lemma 7.1. (*Hanf-locality of first-order logic, see [9]*) *Let Ψ be a first-order sentence with quantifier rank q , and let $r = (3^q - 1)/2$. Let S_1 and S_2 be two structures with a bijection f between their elements such that for every element a of S_1 it holds that the r -neighborhoods of a in S_1 and of $f(a)$ in S_2 are isomorphic. Then S_1 satisfies Ψ iff S_2 satisfies it.*

The proof of Theorem 5.1 is based on the undefinability of graph connectivity by existential monadic second-order sentences [3, 6]. The basic fact is that every existential monadic second-order sentence satisfied by all cycles is also satisfied by some graph which is the union of two cycles. The following lemma is a variant of this result, based on the presentation in [9], with the required modifications. The modifications are needed as a crown is somewhat different from a cycle and due to the presence of bottom elements Hanf-locality cannot be applied directly.

Lemma 7.2. *Let Φ be a universal monadic second order sentence over the language of \leq and equality. If every partial order in \mathcal{R}_1 falsifies Φ then some partial order in \mathcal{R}_0 also falsifies Φ .*

Proof Given an extended crown or extended double crown M , we define a structure G_M over the language E, L_1, L_2, L_3 , where E is a binary relation and L_1, L_2, L_3 are unary relations. The ground sets of M and G_M are the same. Relation E is the comparability relation of the crown or double crown, and L_1, L_2, L_3 correspond to the maximal and minimal elements of the extended crown or extended double crown, resp., the bottom elements. Thus G_M is an undirected graph with vertices colored by the three colors L_1, L_2, L_3 . We refer to G_M as the colored graph of M . The underlying undirected graph for extended crowns (resp., extended double crowns) is complete bipartite graph between a cycle (resp. two cycles) and an independent set. The structures M and G_M are inter-definable by simple first-order sentences. In one direction

$$(a < b) \equiv [E(a, b) \wedge ((L_2(a) \wedge L_1(b)) \vee (L_3(a) \wedge L_2(b)) \vee (L_3(a) \vee L_1(b)))].$$

In the other direction $E(a, b) \equiv (a < b) \vee (b < a)$, $L_3(a) \equiv \min_{\leq}(a)$ and L_2, L_3 can be defined similarly.

For the rest of the argument it is more convenient to switch to existential sentences. We show that if an existential monadic second-order sentence

$$\Phi = \exists A_1, \dots, A_\ell \Psi$$

over the language E, L_1, L_2, L_3 is satisfied by the colored graph G_{M_1} of every extended crown M_1 then it is also satisfied by the colored graph G_{M_2} of some extended double crown M_2 . The lemma then follows directly.

Let q be the quantifier rank of Ψ . Let $r = (3^q - 1)/2$ and $T = 2 \cdot 2^{\ell(2r+1)}$.

Lemma 7.3. *Let M_1 be an extended crown on at least $(4r+4) \cdot T$ elements. For every (A_1, \dots, A_ℓ) -extension G'_{M_1} of G_{M_1} there is an extended double crown M_2 of the same size and an (A_1, \dots, A_ℓ) -extension G'_{M_2} of G_{M_2} such that $G'_{M_1} \equiv_q G'_{M_2}$.*

Proof of Lemma 7.3 Let us consider the substructure C_1 of G_{M_1} corresponding to vertices labeled L_1, L_2 , and its extension (A_1, \dots, A_ℓ) -extension C'_1 . Thus C_1 is a 2-colored cycle. The r -neighborhood of a vertex a in C_1 is the set of vertices which can be reached from a by a path of length at most r . The neighborhood consists of two arcs of r edges each. The r -neighborhoods of vertices labeled L_1 (resp., L_2) in C_1 are isomorphic. The extension adds an additional coloring with 2^ℓ colors. There are T isomorphism types of r -neighborhoods in C'_1 . Hence there are two elements a and b such that their distance on the cycle is at least $2r+2$ and their r -neighborhoods are isomorphic. Let a' and b' be the successors of a and b on the cycle (using some orientation).

Form C'_2 from C'_1 by deleting edges (a, a') and (b, b') , and adding edges (a, b') and (b, a') . It follows from Lemma 7.1 that $C'_1 \equiv_q C'_2$.

Let G'_{M_2} be the extended colored graph obtained from C'_2 by adding colored bottom vertices as in G'_{M_1} . Form G_{M_2} from G'_{M_2} by deleting the unary relations A_1, \dots, A_ℓ and let M_2 be the extended double crown represented by G_{M_2} . The sizes of M_1 and M_2 are the same.

We claim that $G'_{M_1} \equiv_q G'_{M_2}$. It is sufficient to show that Duplicator wins the q -round first-order game on the two structures. As $C'_1 \equiv_q C'_2$, Duplicator wins the q -round first-order game on C'_1 and C'_2 . Also, Duplicator trivially wins the q -round first-order game on the two isomorphic sets of bottom vertices. As all edges are present between the cycles and the bottom vertices, the combination of the two winning strategies gives a winning strategy on G'_{M_1} and G'_{M_2} . $\square \square$

To conclude the proof of Theorem 5.1 assume that \mathcal{R} -revision operators are characterizable by a postulate P . Then \mathcal{R} is universal monadic second-order definable by Lemma 6.5. Thus there is a universal monadic second-order sentence satisfied by every partial order in \mathcal{R}_1 and falsified by every partial order in \mathcal{R}_0 , contradicting Lemma 7.2. \square

References

- [1] C. E. Alchourrón, P. Gärdenfors, and D. Makinson. On the Logic of Theory Change: Partial Meet Contraction and Revision Functions. *J. Symbolic Logic*, 50(2):510–530, 1985.
- [2] H.-D. Ebbinghaus and J. Flum. *Finite Model Theory*. Springer Monographs in Mathematics. Springer-Verlag, Berlin, enlarged edition, 2006.

- [3] R. Fagin. Monadic Generalized Spectra. *Z. Math. Logik Grundlagen Math.*, 21:89–96, 1975.
- [4] E. Fermé and S. O. Hansson. AGM 25 years. Twenty-five Years of Research in Belief Change. *J. Philos. Logic*, 40(2):295–331, 2011.
- [5] A. Grove. Two modellings for theory change. *Journal of Philosophical Logic*, (17):157–170, 1988.
- [6] P. Hájek. On Logics of Discovery. In *MFCS*, pages 30–45, 1975.
- [7] S. O. Hansson. *A Textbook of Belief Dynamics*, volume 11 of *Applied Logic Series*. Kluwer Academic Publishers, Dordrecht, 1999.
- [8] H. Katsuno and A. O. Mendelzon. Propositional Knowledge Base Revision and Minimal Change. *Artificial Intelligence*, 52(3):263–294, 1991.
- [9] L. Libkin. *Elements of Finite Model Theory*. Texts in Theoretical Computer Science. An EATCS Series. Springer-Verlag, Berlin, 2004.
- [10] H. van Ditmarsch, W. van der Hoek, and B. Kooi. *Dynamic Epistemic Logic*. Springer Verlag, 2007.